Private Information Retrieval for Coded Storage

Terence H. Chan, Siu-Wai Ho, and Hirosuke Yamamoto

arXiv:1410.5489v1 [cs.IT] 20 Oct 2014

***Abstract*—Private information retrieval scheme for coded data storage is considered in this paper. We focus on the case where the size of each data record is large and hence only the download cost (but not the upload cost for transmitting retrieval queries) is of interest. We prove that the tradeoff between storage cost and retrieval/download cost depends on the number of data records in the system. We also propose a fairly general class of linear storage codes and retrieval schemes and derive conditions under which our retrieval schemes are error-free and private. Tradeoffs between the storage cost and retrieval costs are also obtained. Finally, we consider special cases when the underlying storage code is based on an MDS code. Using our proposed method, we show that a randomly generated retrieval scheme is indeed very likely to be private and error-free.**

I. Introduction

In modern data storage system, data are usually stored at multiple storage nodes in the cloud. However, system failures are inevitable and in fact common in modern information-technological infrastructure. These failures can be caused by hardware or software failures, or even natural disaster (such as fires, earthquakes and flooding). Another big challenge is how to ensure data stored in the cloud are protected from loss and from being overwritten by illegitimate parties. To ensure that data will not be lost in the case of failures, data redundancy are required. Traditionally, several copies of the same data will be stored. Thus, if one of the copies stored in a node is lost, the lost content can be recovered by retrieving a copy of the data from other surviving nodes.

Such replication approach is not efficient in terms of storage cost. An alternative approach is to use erasure codes (such as Reed-Solomon codes, an example of maximum distance separable codes) to store the data. Thus, any content loss can then be seen as a data erasure and be recovered from the code. Comparing with replication, this coding approach has a lower storage cost but also a higher repair cost (measured by the amount of data sent to repair a failure node). There is a fundamental tradeoff between the costs for storage and repair. In [1], regenerating codes were proposed to achieve the optimal tradeoff between storage cost and repair cost. This approach has since received a lot of attentions [2]–[5].

In some data storage applications, not only the data protection against system failure and illegitimate parties is a concern but also the privacy of data retrieval query needs to be protected. For example, consider investors retrieving stock prices from a cloud storage provider (CSP). In this case, it is preferred to keep the identity of the stocks secret from other users and the CSP. Private information retrieval (PIR) ensures that servers storing data would not know the identity of the data record that a user is retrieving. One naive approach to achieve PIR is by downloading every record in the system no matter which record the user is needed. The downside of this approach is the extremely large retrieval cost, which linearly increases with N (the number of records stored).

The challenge however is how to design an efficient PIR scheme (in terms of costs for data storage and retrieval). In a k-server PIR scheme, a user sends k queries to k servers which will reply to the user k answers. The retrieval efficiency of a PIR scheme is measured by the upload cost (i.e., the size of the queries) and the retrieval cost (i.e., the size of the answers). The sum of these two costs is referred to as the total communication complexity. In [6], two very interesting schemes were proposed. The first scheme achieves total communication complexity with $O(N^{\frac{1}{K}})$ where K is the number of servers and N is the number of data records. If there are $\frac{1}{3}\log_2(N)+1$ servers, the second scheme achieves a total communication complexity with $\frac{1}{3}(1+o(1))\cdot\log_2^2(N)\cdot log_2(log_2(2N))$. These results were further elaborated in [7, Sec. 3.1-3.4]. For data records of size equal to ℓ, a PIR scheme with total communication complexity equal to 4ℓ was proposed. Note that all these PIR schemes require that each storage node stores a copy of every data records. This paper considers a more general scenario where this requirement is not necessarily satisfied.

In [8], a single-database "computational" PIR scheme to achieve retrieval cost less than N was proposed. Unlike previous information-theoretic schemes, the privacy of the identity of the retrieved data is protected based on the computational hardness of the quadratic residuosity problem. This scheme can achieve N^{ϵ} for any $\epsilon > 0$ by assuming the hardness of deciding quadratic residuosity. In [9], [10], PIR schemes which achieve a retrieval cost of $\log^2 N$ were proposed. However, the computational complexity of these schemes is of concern. In fact, a scheme which requires less computation of public-key operations was shown in [11]. Information-theoretic approach has also been considered to achieve PIR [12]. A more general model was considered in [13] in which 1) not every server may respond to the query and 2) some servers are tampered and hence erroneous responses will be returned to the users. There are many other works which considered similar PIR problems. Due to limited space, please refer to [7] for more details.

Almost all of the existing works on PIR assume uncoded data storage (where each storage node stores a copy of every data record). However, uncoded storage (or more precisely, storage based on repetition codes) has the highest storage cost. To reduce the storage cost, coded storage (e.g., ones that base on maximum distance separable (MDS) codes, or locally repairable codes) have been proposed. Motivated by

the advances in coded data storage, this paper considers PIR schemes for coded data storage system.

To the best of our knowledge, the recent work [14] is the only existing work that discusses private information retrieval in coded storage. However, there are some fundamental differences between the approach in [14] and the one considered in this paper. In [14], it was proved that only one extra bit is required to download to ensure the privacy of the identity of retrieved record. However, the extremely low retrieval cost comes with a price that the number of storage nodes (and also the storage cost) needs to grow with the data record size. We however assume that the number of storage nodes is fixed.

In [14, Section V], PIR schemes for MDS codes based storage were proposed. The optimality of these schemes were not discussed. In contrast, this paper considers a fairly general framework where we impose no restriction on the choice of the underlying storage codes. We obtain the tradeoffs between storage cost and retrieval cost and show that the optimal tradeoff can be achieved if MDS codes are used. Unlike in [14] where the number of data records is assumed to be fixed, we will also show that the optimal tradeoff between the storage and retrieval costs also depend on the number of records in the system.

The organisation of the paper is as follows. Section II is the background for storage codes and private information retrieval (PIR) scheme. In Section III, we propose our PIR schemes for coded storage. Sufficient conditions for which our scheme is private will be derived. In Section IV, we will identify the tradeoff between storage cost and retrieval cost.

Notations: For random variables X, Y, the entropy of X will be denoted by $H(X)$, while the mutual information between X and Y is denoted by $I(X;Y)$. A set of random variables $(X_i, i \in \alpha)$ can also be denoted by X_α. Random variables are usually denoted by capital letters (such as X, Y etc) while their realisations by small letters. Sample spaces over which random variables are defined will be denoted by calligraphic letters.

II. Problem Formulation

A. System Formulation

We assume that there are N data records $(D_n,\ n \in \mathcal{N} \triangleq \{1, \ldots, N\})$ to be stored in K data storage nodes (indexed by $\mathcal{K} \triangleq \{1, \ldots, K\}$). Assume without loss of generality that these data records are independently and uniformly distributed over a common sample space $\mathcal{D}$. In other words,

$$\Pr(D_n = d_n, n \in \mathcal{N}) = 1/|\mathcal{D}|^N. \tag{1}$$

Hence, $H(D_\mathcal{N}) = \sum_{n\in\mathcal{N}} H(D_n)$ and $H(D_n) = \log|\mathcal{D}|$.

Definition 1 (Storage Code): For any $k \in \mathcal{K}$, let X_k (defined over $\mathcal{X}_K$) be the content stored in storage node k. A (N, K) *storage code* is defined by an encoding mapping

$$\Omega:\ \mathcal{D}^N \to \prod_{k\in\mathcal{K}} \mathcal{X}_k$$

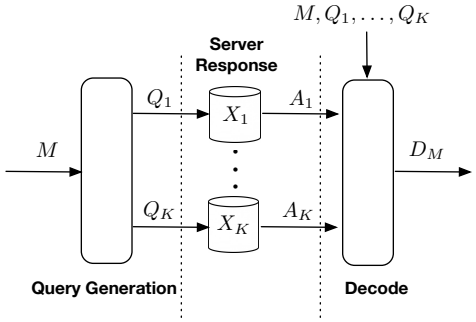

Figure 1. Retrieving data record D_M

such that $(X_1, \ldots, X_K) = \Omega(D_\mathcal{N})$.

Clearly, it is required that all data $D_1, \ldots, D_N$ can be retrieved from the contents in the K storage nodes. In other words, $H(D_1, \ldots, D_N | X_1, \ldots, X_K) = 0$. In addition, there are often other criteria that a storage code needs to satisfy. For instance, contents stored at a storage node may be corrupted (due to software or hardware failures) and hence need to be recovered from other storage nodes. For example, one may require that for any $k \in \mathcal{K}$, there exists a subset $\alpha \subseteq \mathcal{K} \setminus \{k\}$ of size at most r such that $H(X_k | X_j, j \in \alpha) = 0$ [3], [15]. In other scenarios, it may also be required that all data can still be recovered when there are multiple (no more than δ) simultaneous node failures. In this case, it is required that

$$H(X_\mathcal{K} | X_j, j \in \beta) = 0$$

for any $\beta \subseteq \mathcal{K}$ of size $K - \delta$.

For any storage code Ω, we will use $\mathsf{SC}(\Omega)$ to denote the *storage cost*, which is the number of bits each storage node must use to store one information bit. More precisely,

$$\mathsf{SC}(\Omega) \triangleq \max_{k\in\mathcal{K}} \frac{\log|\mathcal{X}_k|}{N \log|\mathcal{D}|}. \tag{2}$$

Remark 1: We do not assume any compression schemes used in the storage and hence $\log|\mathcal{X}_k|$ is used instead of $H(X_k)$. As $\log|\mathcal{X}_k| \geq H(X_k)$, results obtained in this paper will still hold if $H(X_k)$ is used. Also, there are other alternative storage cost measures (e.g., "the total storage cost" $\sum_{k\in\mathcal{K}} \log|\mathcal{X}_k| / N\log|\mathcal{D}|$). We consider (2) to focus on the required storage cost on individual nodes. With respect to our code (to be defined in Section III), the two measures differed only by a constant multiplicative factor K (which is the number of storage nodes and is fixed).

Clearly, a data storage will be useless if the stored data cannot be retrieved. A *retrieval scheme* usually consists of three steps (see Figure 1). The first step is the *queries generation*. Let $M \in \mathcal{N}$ be the index of the data record to be retrieved. We will assume that M is uniformly distributed over $\mathcal{N}$. To retrieve D_M, a user will first generate K queries $(Q_k, k \in \mathcal{K})$

according to a conditional probability distribution[1]

$$\Lambda_1(q_k, k \in \mathcal{K}|m) = \Pr(Q_k = q_k, k \in \mathcal{K}|M = m). \quad (3)$$

Then, the query Q_k will be sent to storage node k. We call Λ_1 the "*query generation distribution*".

In the second step, each storage node k, receiving Q_k, will generate a response via the function

$$\Lambda_{2,k} : \mathcal{X}_k \times \mathcal{Q}_k \to \mathcal{A}_k. \quad (4)$$

In particular, $A_k \triangleq \Lambda_{2,k}(X_k, Q_k)$ will be computed and returned to the user. The mappings $(\Lambda_{2,k}, k \in \mathcal{K})$ are called "*response mappings*". Finally, in the third step, the user will "*decode*" D_M from $(M, Q_k, A_k, k \in \mathcal{K})$.

Definition 2: A retrieval scheme Λ is specified by the query generation distribution and the response mappings

$$\Lambda \triangleq (\Lambda_1, \Lambda_{2,k}, k \in \mathcal{K}).$$

Remark 2: Once the storage code Ω and the retrieval scheme Λ are given, the probability distribution of $(M, X_k, Q_k, A_k, k \in \mathcal{K}, D_n, n \in \mathcal{N})$ is well-defined. In particular,

$$(X_{\mathcal{K}}, D_{\mathcal{N}}) \perp (M, Q_{\mathcal{K}}) \quad (5)$$

$$H(A_k|Q_k, X_k) = 0 \quad (6)$$

$$H(D_n, n \in \mathcal{N}) = \sum_{n \in \mathcal{N}} H(D_n). \quad (7)$$

Conversely, if the probability distribution of $(X_k, k \in \mathcal{K}, D_n, n \in \mathcal{N})$ is known, then one can also deduce the underlying storage code Ω. Similarly, the probability distributions of $(M, Q_k, k \in \mathcal{K})$ (and (Q_k, A_k, X_k)) respectively specify Λ_1 (and $\Lambda_{2,k}$). Therefore, in this paper, we often refer a storage system by the set of random variables $(M, X_k, Q_k, A_k, k \in \mathcal{K}, D_n, n \in \mathcal{N})$.

Definition 3 (Error-free Retrieval): For a given storage code Ω, a retrieval scheme Λ is called *error-free* if

$$H(D_M|M, Q_{\mathcal{K}}, A_{\mathcal{K}}) = 0. \quad (8)$$

Definition 4 (Private Retrieval): Let Φ be a collection of subsets of $\mathcal{K}$. A retrieval scheme Λ is called Φ-*private* if

$$I(M; Q_j, j \in \alpha) = 0 \quad (9)$$

for all subset $\alpha \in \Phi$. We call α the *collusion pattern*.

According to our definitions, if a retrieval scheme is error-free and Φ-*private*, then 1) a user can always retrieve (and decode) its interested data record D_M with no errors, and 2) for any $\alpha \in \Phi$, the colluding set of storage nodes (i.e., indexed by α) obtains no information about the identity (the index M) but not necessarily the content of the record being retrieved (i.e., D_M).

[1] In this paper, we consider the general case that queries are randomly generated to prevent the storage nodes from knowing the identity of the data record to be retrieved.

The *retrieval cost* $\mathsf{RC}(\Lambda, \Omega)$ of a retrieval scheme Λ (with respect to an underlying storage code Ω) is defined by

$$\mathsf{RC}(\Lambda, \Omega) \triangleq \max_{k \in \mathcal{K}} \frac{\log |\mathcal{A}_k|}{\log |\mathcal{D}|}.$$

Since each D_n is uniformly distributed over $\mathcal{D}$,

$$\mathsf{RC}(\Lambda, \Omega) = \max_{k \in \mathcal{K}} \frac{\log |\mathcal{A}_k|}{H(D_1)} \geq \max_{k \in \mathcal{K}} \frac{H(A_k)}{H(D_1)}.$$

The retrieval cost $\mathsf{RC}(\Lambda, \Omega)$ measures the number of bits required to retrieve from a server node, for each information bit of a record.

Remark 3: In this paper, we consider the case when the size of each data record is large. Due to (5) and (8), $H(A_{\mathcal{K}}) \geq H(D_M)$. Thus, the retrieval cost must grow with the size of data. However, this is not necessary the case for upload cost (transmission bandwidths required to send the queries to the storage nodes). To illustrate, suppose we are given a (N, K) storage code Ω, together with a retrieval scheme Λ, such that each data record is of size $\log |\mathcal{D}|$ bits. Now, suppose the size of each data record is in fact $c \log |\mathcal{D}|$ bits. In this case, we can partition each data record into c chunks, each of size $\log |\mathcal{D}|$ bits. For $i = 1, \ldots, c$, we may store the i^{th} chunk of all the data records using the storage code Ω. To retrieve a data record, a user can retrieve all the c chunks of the record using the retrieval scheme Λ. The catch however is that the retrieval query sent to a storage node can be the same for every chunk. Hence, only one (not c) query is sent, and the upload cost stays the same for all c. In addition, similar to the case of storage cost, we use $\log |\mathcal{A}_k|$ instead of $H(A_{\mathcal{K}})$. Yet, results obtained in this paper will still hold if $H(A_{\mathcal{K}})$ is used.

B. Achievability

A core question in private information retrieval is to determine the fundamental tradeoffs between storage and retrieval costs of error-free and private retrieval schemes. To facilitate our discussion, we will need the following definitions.

Definition 5 (Achievability): With respect to a given K and collusion patterns Φ, a *storage-retrieval cost tuple* $(\mathsf{SC}^*, \mathsf{RC}^*)$ is called N-achievable if there exists a (N, K) storage code Ω and a Φ-private error free retrieval scheme Λ such that

$$\mathsf{SC}(\Omega) \leq \mathsf{SC}^* \quad (10)$$

$$\mathsf{RC}(\Lambda, \Omega) \leq \mathsf{RC}^*. \quad (11)$$

Furthermore, it is called *strongly achievable* if there exists a sequence (indexed by a) of $N^{(a)}$-achievable tuples $(\mathsf{SC}^{(a)}, \mathsf{RC}^{(a)})$ such that

1) $\lim_{a \to \infty} N^{(a)} = \infty$
2) $\limsup_{a \to \infty} \mathsf{SC}^{(a)} \leq \mathsf{SC}^*$
3) $\limsup_{a \to \infty} \mathsf{RC}^{(a)} \leq \mathsf{RC}^*$.

Remark 4: Clearly, a strongly achievable storage-retrieval cost tuple is also N-achievable (for any fixed N). However, as we shall prove in Theorem 1, the converse is not true in general.

Example 1: Let $D_1, \ldots, D_N$ be the data records where $D_n \in GF(q)$ for $1 \leq n \leq N$. Construct a PIR scheme as follows: Let Q_0 be the vector whose M^{th} entry is one and all the others are equal to zero. Encode the "secret message" Q_0 with a linear secret sharing scheme that is defined by K random shares $Q_1, \ldots, Q_K$ such that

1) for any collusion pattern $\alpha \in \Phi$, $I(Q_0; Q_k, k \in \alpha) = 0$
2) there exists coefficients $c_1, \ldots, c_K$ such that

$$Q_0 = \sum_{k=1}^{K} c_k Q_k. \tag{12}$$

Then, Q_k will be sent to storage node k for $k \in \mathcal{K}$. Upon receiving Q_k, node k will return the response $A_k \triangleq [D_1, \ldots, D_N]Q_k$. By (12), we have

$$D_M = [D_1, \ldots, D_N]Q_0 \tag{13}$$

$$= [D_1, \ldots, D_N] \sum_{k=1}^{K} c_k Q_k \tag{14}$$

$$= \sum_{k=1}^{K} c_k A_k. \tag{15}$$

Therefore, the user can decode the desired data record D_M. Furthermore, by the definition of the secret sharing code, $I(M; Q_j) = 0$ for all $\alpha \in \Phi$. Hence, the so-constructed scheme is indeed a Φ-private error-free retrieval scheme.

The above example requires that each storage node stores every data record. Therefore, the storage cost is the highest (equal to 1). The following example however considers coded storage.

Example 2 (PIR for coded storage): We consider $N = 2$. Each of the two record contains two bits. Assume without loss of generality that the first record is (a_1, b_1) and the second one is (a_2, b_2). There are $K = 3$ storage nodes, each of which will store two bits. The first server will store $X_1 = (a_1, b_1)$, the second $X_2 = (a_2, b_2)$ and the third $X_3 = (a_1 \oplus a_2, a_2 \oplus b_2)$. Clearly, the storage cost of this code is only $1/2$.

	server 1	server 2	server 3
$Q_k = 1$	a_1	a_2	$a_1 \oplus a_2$
$Q_k = 2$	b_1	b_2	$b_1 \oplus b_2$
$Q_k = 3$	$a_1 \oplus b_1$	$a_2 \oplus b_2$	$a_1 \oplus a_2 \oplus b_1 \oplus b_2$

Table I
RESPONSE MAPPINGS

Let Φ be the set of all α such that $|\alpha| = 1$. We will assume that each query Q_k (for $k = 1, 2, 3$) is a ternary random variables taking values from the set $\{1, 2, 3\}$. Table I defines how A_k is generated from Q_k and X_k. For example, if $Q_2 = 2$, then A_2 (the response generated by server 2) will be equal to b_2. If $Q_3 = 3$, then $A_3 = a_1 \oplus a_2 \oplus b_1 \oplus b_2$. Now, we will define the conditional probability distribution of (Q_1, Q_2, Q_3) given M (the index of the record to be retrieved) such that

$$\Pr(Q_1 = q_1, Q_2 = q_2, Q_3 = q_3 | M = 1) = 1/3$$

if $(q_1, q_2, q_3) \in \{(1, 3, 3), (2, 1, 1), (3, 2, 2)\}$, and 0 otherwise. Similarly, we define

$$\Pr(Q_1 = q_1, Q_2 = q_2, Q_3 = q_3 | M = 2) = 1/3$$

if $(q_1, q_2, q_3) \in \{(3, 1, 3), (1, 2, 1), (2, 3, 2)\}$. It can be verified directly that

$$I(Q_k; M) = 0$$

for all $k = 1, 2, 3$ and hence the retrieval scheme is private. In addition, the user can always reconstruct the desired record. For example, when $M = 1$ and

$$(q_1, q_2, q_3) = (1, 3, 3),$$

then the user will receive a_1, $a_2 \oplus b_2$ and $a_1 \oplus b_1 \oplus a_2 \oplus b_2$ from the three servers. Clearly, in this case, the user can decode a_1 and b_1 back. The retrieval cost of this scheme is equal to $1/2$. Hence, $(1/2, 1/2)$ is 2-achievable.

Theorem 1: Suppose $K = 3$ and

$$\Phi = \{\{k\} : k \in \mathcal{K}\}.$$

Then the storage-retrieval cost tuple $(1/2, 1/2)$ is 2-achievable but not strongly achievable.

Proof: The 2-achievability of the storage-retrieval cost tuple $(1/2, 1/2)$ follows from Example 2. The proof for that it is not strongly achievable will be given in Appendix A. ■

As a corollary from Theorem 1, N-achievability and strong achievability are in fact different. This implies that the tradeoff between the storage cost and retrieval cost does depend on the number of data records in the system. To our best knowledge, this is the first proof for such a phenomenon.

III. LINEAR PIR CODES

Example 1 is a PIR scheme constructed from a secret sharing scheme. It requires that each storage node contains all the N records and hence its storage cost is the highest (equal to 1). In contrast, the storage cost of the code in Example 2 is only 1/2. In this section, we will propose a class of storage codes, together with its corresponding retrieval scheme. As we shall see in Section IV, our class of codes can be constructed over a wide range of storage-retrieval cost tuples.

We will first begin with the description of the storage code. In our proposed scheme, we denote the N data records by $\mathbf{d}_1, \ldots, \mathbf{d}_N$ each of which is a vector of length $(K - S)L$ over $GF(q)$ for some positive integer constants L and S (whose physical meaning will become clear later). In our proposed scheme, each data record will be "encoded and stored" separately. Specifically, there exists K mappings

$$g_k : \mathcal{D} \to GF(q)^L, \quad k \in \mathcal{K}$$

such that $g_k(\mathbf{d}_n)$ is the "coded data piece" of $\mathbf{d}_n$ being stored in storage node k. Therefore, for data records $\mathbf{d}_1, \ldots, \mathbf{d}_N$, the storage node k will store the length LN vector

$$X_k = [g_k(\mathbf{d}_1), \ldots, g_k(\mathbf{d}_N)]^\top$$

$$\mathbf{V} = \left[\begin{array}{ccc|ccc|ccc|ccc} v^*_{1,1,1} & \cdots & v^*_{1,R,1} & v^*_{1,1,2} & \cdots & v^*_{1,R,2} & \cdots & \cdots & v^*_{1,1,K} & \cdots & v^*_{1,R,K} \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \ddots & \ddots & \vdots & \ddots & \vdots \\ v^*_{L,1,1} & \cdots & v^*_{L,R,1} & v^*_{L,1,2} & \cdots & v^*_{L,R,2} & \cdots & \cdots & v^*_{L,1,K} & \cdots & v^*_{L,R,K} \\ v_{1,1,1} & \cdots & v_{1,R,1} & v_{1,1,2} & \cdots & v_{1,R,2} & \cdots & \cdots & v_{1,1,K} & \cdots & v_{1,R,K} \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \ddots & \ddots & \vdots & \ddots & \vdots \\ v_{T,1,1} & \cdots & v_{T,R,1} & v_{T,1,2} & \cdots & v_{T,R,2} & \cdots & \cdots & v_{T,1,K} & \cdots & v_{T,R,K} \end{array}\right] \tag{16}$$

In order to retrieve $\mathbf{d}_n$ without errors, it is necessary that the data record $\mathbf{d}_n$ and $[g_k(\mathbf{d}_n)^\top, k \in \mathcal{K}]$ are one-to-one corresponded. Therefore, $\mathbf{d}_n$ and $[g_k(\mathbf{d}_n)^\top, k \in \mathcal{K}]$ can be used interchangeably. In fact, for notation simplicity, we will simply write $\mathbf{d}_n$ as

$$[g_k(\mathbf{d}_n)^\top, k \in \mathcal{K}] = \left[\begin{array}{ccc} d_{n,1,1} & \cdots & d_{n,1,K} \\ \vdots & \ddots & \vdots \\ d_{n,L,1} & \cdots & d_{n,L,K} \end{array}\right] \tag{17}$$

such that

$$g_k(\mathbf{d}_n) = [d_{n,1,k}, \ldots, d_{n,L,k}]$$

is the transpose of the k^{th} column of $\mathbf{d}_n$.

Now, to define the storage code, it suffices to define g_k, or equivalently, the condition that $[g_k(\mathbf{d}_n)^\top, k \in \mathcal{K}]$ must satisfy. In this paper, we will consider a wide class of linear storage codes such that

$$\mathbf{d}_n \mathbf{P} = \mathbf{0} \tag{18}$$

for a full-rank $K \times S$ parity check matrix

$$\mathbf{P} = \left[\begin{array}{ccc} p_{1,1} & \cdots & p_{1,S} \\ \vdots & \ddots & \vdots \\ p_{K,1} & \cdots & p_{K,S} \end{array}\right]. \tag{19}$$

The choice of $\mathbf{P}$ determines the type of storage codes that are being used. For uncoded storage, $\mathbf{P}$ is the $K \times K-1$ full rank matrix such that

$$[1, 1, \ldots, 1]\mathbf{P} = \mathbf{0}.$$

Following our convention, the storage node k will store a length LN column vector

$$X_k \triangleq [d_{1,1,k}, \ldots, d_{1,L,k}, d_{2,1,k}, \ldots, d_{2,L,k}, \ldots, d_{N,L,k}]^\top,$$

which is obtained by concatenating the k^{th} column of each record. As a result, it is obvious that

$$\sum_{k=1}^{K} p_{k,s} X_k = \mathbf{0}. \tag{20}$$

For notation simplicity, we will refer our code as $\Delta_{\mathbf{P},N,L}$ (or just $\Delta_{\mathbf{P}}$ when N and L are understood implicitly). It can be directly verified that its storage cost is

$$\mathsf{SC}(\Delta_{\mathbf{P}}) = 1/(K-S). \tag{21}$$

Next, we will define a general class of information retrieval schemes for the above storage code $\Delta_{\mathbf{P}}$. Our proposed retrieval scheme is defined by a $(T+L) \times RK$ matrix $\mathbf{V}$ (see (16)), whose columns will be respectively labeled by

$$V_{1,1}, \ldots, V_{R,1}, V_{1,2}, \ldots, V_{R,2}, \ldots, V_{1,K}, \ldots, V_{R,K}.$$

We will refer the retrieval scheme as $\Theta_{\mathbf{V}}$. A user can retrieve the Mth record $\mathbf{d}_M$ by following the steps below:

1) (Initialisation) Generate a $LN \times T$ random matrix $\mathbf{U}$, whose entries are all randomly and independently selected from $GF(q)$. Let the columns of $\mathbf{U}$ be respectively
$$U_t, t \in \mathcal{T} \triangleq \{1 \ldots, U_T\}.$$
2) (Query Generation) For each $k \in \mathcal{K}$, the user will generate R "query" vectors $Q_{1,k} \ldots Q_{R,k}$ such that
$$Q_{r,k} \triangleq \sum_{t=1}^{T} v_{t,r,k} U_t + \sum_{\ell=1}^{L} v^*_{\ell,r,k} E_{M,\ell}. \tag{22}$$
Here, $v_{t,r,k}$ and $v^*_{\ell,r,k}$ are defined as in (16) and $E_{M,\ell}$ is a length LN column vector whose entry is 1 on the $L(M-1)+\ell$ th row and is 0 otherwise. In other words, $E_{M,\ell}$ is the vector such that
$${E_{M,\ell}}^\top X_k = d_{M,\ell,k} \tag{23}$$
for all $M \in \mathcal{N}, \ell \in \mathcal{L} \triangleq \{1, \ldots, L\}$ and $k \in \mathcal{K}$.
3) (Response Mappings) The collection of the R query vectors
$$(Q_{r,k}, r \in \mathcal{R} \triangleq \{1, \ldots, R\})$$
will then be sent to storage node k, which will respond by computing and returning to user the tuple $(A_{r,k}, r = 1, \ldots, R)$ where $A_{r,k} = Q_{r,k}^\top X_k$. Thus,
$$A_{r,k} = \sum_{t=1}^{T} v_{t,r,k} U_t^\top X_k + \sum_{\ell=1}^{L} v^*_{r,\ell,k} d_{M,\ell,k}. \tag{24}$$
As each storage node will return R symbols back to the user, the retrieval cost is thus
$$\mathsf{RC}(\Theta_{\mathbf{V}}, \Delta_{\mathbf{P}}) = \frac{R}{L(K-S)}. \tag{25}$$

Ideally, the data record

$$\mathbf{d}_M = (d_{M,\ell,k}, \ell \in \mathcal{L}, k \in \mathcal{K})$$

should be obtained by solving the system of linear equations

(24), together with

$$U_t^\top \sum_{k=1}^{K} p_{k,s} X_k = 0, \tag{26}$$

$$\sum_{k=1}^{K} p_{k,s} d_{M,\ell,k} = 0, \tag{27}$$

for all $s \in \mathcal{S} \triangleq \{1, \ldots, S\}$ and $\ell \in \mathcal{L}$. However, a retrieval scheme $\Theta_{\mathbf{V}}$ is not necessarily error-free (where $\mathbf{d}_M$ can be solved without errors) or private. The following theorem gives conditions on $\mathbf{V}$ (with respect to the storage code $\Delta_{\mathbf{P}}$) under which our retrieval scheme is in fact error-free and private.

Theorem 2: The retrieval scheme $\Theta_{\mathbf{V}}$ is error-free if the following system of linear equations

$$\begin{cases} \sum_{k=1}^{K} w_{t,k} p_{k,s} = 0, & \forall t \in \mathcal{T}, s \in \mathcal{S} \\ \sum_{k=1}^{K} w^*_{\ell,k} p_{k,s} = 0, & \forall \ell \in \mathcal{L}, s \in \mathcal{S} \\ \sum_{t=1}^{T} v_{t,r,k} w_{t,k} + \sum_{\ell=1}^{L} v^*_{\ell,r,k} w^*_{\ell,k} = 0, & \forall k \in \mathcal{K}, r \in \mathcal{R} \end{cases} \tag{28}$$

has a unique solution, where the unknowns are

$$(w_{t,k}, t \in \mathcal{T}, k \in \mathcal{K}, w^*_{l,k}, l \in \mathcal{L}, k \in \mathcal{K}).$$

The condition is called *retrievability condition*.

Furthermore, the retrieval scheme is Φ-private if

$$\langle V_{r,k}, r \in \mathcal{R}, k \in \alpha \rangle \cap V_0 = \{\mathbf{0}\} \tag{29}$$

for any collusion pattern $\alpha \in \Phi$. Here, V_0 is the L dimensional vector space containing all vectors of the form $[c_1, \ldots, c_L, 0, \ldots, 0]^\top$. We call (29) the *privacy condition*.

Proof: For any $t \in \mathcal{T}$, $r \in \mathcal{R}$ and $k \in \mathcal{K}$, let

$$w_{t,k} \triangleq U_t^\top X_k \tag{30}$$

$$w^*_{\ell,k} \triangleq d_{M,\ell,k}. \tag{31}$$

Then rewriting (24), we have

$$A_{r,k} = \sum_{t=1}^{T} v_{t,r,k} w_{t,k} + \sum_{\ell=1}^{L} v^*_{r,\ell,k} w^*_{\ell,k}. \tag{32}$$

Invoking (20), for all $t \in \mathcal{T}$ and $s \in \mathcal{S}$, we have

$$U_t^\top \left(\sum_{k=1}^{K} p_{k,s} X_k \right) = 0. \tag{33}$$

Consequently,

$$\sum_{k=1}^{K} p_{k,s} w_{t,k} = 0. \tag{34}$$

Similarly, by (18), we have

$$\sum_{k=1}^{K} p_{k,s} w^*_{\ell,k} = 0, \quad \forall \ell \in \mathcal{L}, s \in \mathcal{S}. \tag{35}$$

As the matrices $\mathbf{P}$ and $\mathbf{V}$ are chosen to satisfy the retrievability condition (that (28) has a unique solution), the system of linear equations (32), (34) and (35) also has a unique solution. Hence, the user can uniquely decode $\mathbf{d}_M$. Therefore, D_M is a function of $M, Q_{\mathcal{K}}, A_{\mathcal{K}}$ and hence $\Theta_{\mathbf{V}}$ is error-free.

Next, we will prove that $\Theta_{\mathbf{V}}$ is Φ-private if (29) is satisfied. Let $\alpha \in \Phi$ be a collusion pattern. Assume without loss of generality that $\alpha = \{1, \ldots, |\alpha|\}$. Let $\mathbf{G}$ and $\mathbf{G}^-$ be two submatrices of $\mathbf{V}$ as defined in (36) and (37). Note that both $\mathbf{G}$ and $\mathbf{G}^-$ have exactly $|\alpha|R$ columns. By invoking elementary column operations, we can prove the existence of an invertible matrix $\mathbf{C}$ such that $\mathbf{G}^-\mathbf{C}$ has rank $\delta \leq T$ and

$$\mathbf{G}^-\mathbf{C} = \begin{bmatrix} f_{1,1} & f_{1,2} & \cdots & f_{1,\delta} & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ f_{T,1} & f_{T,2} & \cdots & f_{T,\delta} & 0 & \cdots & 0 \end{bmatrix}$$

is in the reduced column echelon form. In particular, for each $a = 1, \ldots, \delta$, there exists an index $1 \leq t_a \leq T$ such that

$$f_{t_a,a} \neq 0 \tag{38}$$

$$f_{t_a,b} = 0, \quad \forall b \neq a. \tag{39}$$

Now, consider the matrix

$$\mathbf{GC} = \begin{bmatrix} f^*_{1,1} & f^*_{1,2} & \cdots & f^*_{1,\delta} & f^*_{1,\delta+1} & \cdots & f^*_{1,R|\alpha|} \\ \vdots & \vdots & \ddots & \vdots & \vdots & \cdots & \vdots \\ f^*_{L,1} & f^*_{L,2} & \cdots & f^*_{L,\delta} & f^*_{L,\delta+1} & \cdots & f^*_{L,R|\alpha|} \\ f_{1,1} & f_{1,2} & \cdots & f_{1,\delta} & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ f_{T,1} & f_{T,2} & \cdots & f_{T,\delta} & 0 & \cdots & 0 \end{bmatrix}.$$

The privacy condition (29) guarantees that

$$f^*_{\ell,\delta+1} = \cdots = f^*_{\ell,R|\alpha|} = 0$$

for all $\ell = 1, \ldots, L$. Now, let

$$\left[Q^*_1, \ldots, Q^*_{R|\alpha|}\right] = \left[Q_{1,1}, \ldots, Q_{R,|\alpha|}\right] \mathbf{C} \tag{40}$$

and $\mathbf{U}^* = [E_{M,1}, \cdots, E_{M,L}, \mathbf{U}]$. Then it can be verified directly that

$$\left[Q_{1,1}, \ldots, Q_{R,|\alpha|}\right] = \mathbf{U}^*\mathbf{G}.$$

Hence,

$$\begin{aligned} &[Q^*_1, \ldots, Q^*_{R|\alpha|}] \\ &\quad = [Q_{1,1}, \ldots, Q_{R,|\alpha|}]\mathbf{C} \\ &\quad = \mathbf{U}^*\mathbf{G}\mathbf{C} \\ &\quad = \mathbf{U}^* \begin{bmatrix} f^*_{1,1} & f^*_{1,2} & \cdots & f^*_{1,\delta} & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \cdots & \vdots \\ f^*_{L,1} & f^*_{L,2} & \cdots & f^*_{L,\delta} & 0 & \cdots & 0 \\ f_{1,1} & f_{1,2} & \cdots & f_{1,\delta} & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ f_{T,1} & f_{T,2} & \cdots & f_{T,\delta} & 0 & \cdots & 0 \end{bmatrix}. \end{aligned}$$

Note that $Q^*_i = 0$ for all $i = \delta + 1, \ldots, R|\alpha|$. Now, for any

$$\mathbf{G} = \left[\begin{array}{ccc|ccc|ccc|ccc} v^*_{1,1,1} & \cdots & v^*_{1,R,1} & v^*_{1,1,2} & \cdots & v^*_{1,R,2} & \cdots & \cdots & v^*_{1,1,|\alpha|} & \cdots & v^*_{1,R,|\alpha|} \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \ddots & \ddots & \vdots & \ddots & \vdots \\ v^*_{L,1,1} & \cdots & v^*_{L,R,1} & v^*_{L,1,2} & \cdots & v^*_{L,R,2} & \cdots & \cdots & v^*_{L,1,|\alpha|} & \cdots & v^*_{L,R,|\alpha|} \\ v_{1,1,1} & \cdots & v_{1,R,1} & v_{1,1,2} & \cdots & v_{1,R,2} & \cdots & \cdots & v_{1,1,|\alpha|} & \cdots & v_{1,R,|\alpha|} \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \ddots & \ddots & \vdots & \ddots & \vdots \\ v_{T,1,1} & \cdots & v_{T,R,1} & v_{T,1,2} & \cdots & v_{T,R,2} & \cdots & \cdots & v_{T,1,|\alpha|} & \cdots & v_{T,R,|\alpha|} \end{array}\right] \tag{36}$$

$$\mathbf{G}^- = \left[\begin{array}{ccc|ccc|ccc|ccc} v_{1,1,1} & \cdots & v_{1,R,1} & v_{1,1,2} & \cdots & v_{1,R,2} & \cdots & \cdots & v_{1,1,|\alpha|} & \cdots & v_{1,R,|\alpha|} \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \ddots & \ddots & \vdots & \ddots & \vdots \\ v_{T,1,1} & \cdots & v_{T,R,1} & v_{T,1,2} & \cdots & v_{T,R,2} & \cdots & \cdots & v_{T,1,|\alpha|} & \cdots & v_{T,R,|\alpha|} \end{array}\right] \tag{37}$$

$a = 1, \ldots, \delta$,

$$Q^*_a = \sum_{t=1}^{T} f_{t,a} U_t + \sum_{\ell=1}^{L} f^*_{\ell,a} E_{M,\ell} \tag{41}$$

$$= f_{t_a,a} U_{t_a} + \sum_{t \neq t_a} f_{t,k} U_t + \sum_{\ell=1}^{L} f^*_{\ell,a} E_{M,\ell}. \tag{42}$$

Clearly, each Q^*_a can be written as a sum of columns in $\mathbf{U}^*$. However, by (38)-(39), the column U_{t_a} appears only in the sum for Q^*_a. As entries in each U_{t_a} are uniformly and independently distributed $GF(q)$, we can in fact prove that each Q^*_a in $(Q^*_1, \ldots, Q^*_{R|\alpha|})$ is uniformly and independently distributed to each other, no matter what M is. In other words,

$$I(Q^*_1, \ldots, Q^*_{R|\alpha|}; M) = 0. \tag{43}$$

Finally, by (40), we have

$$I(Q_{1,1}, \ldots, Q_{R,|\alpha|}; M) = 0. \tag{44}$$

As (44) holds for all $\alpha \in \Phi$, our retrieval scheme is indeed Φ-private. And the theorem is proved. ∎

Remark 5: It is worth to mention that our codes and retrieval schemes are well-defined irrespective of N. Therefore, if a storage-retrieval cost tuple $(\mathsf{SC}, \mathsf{RC})$ can be achieved using our proposed schemes, the tuple is also strongly achievable.

IV. Tradeoff analysis

In the previous section, we proposed a linear storage code $\Delta_{\mathbf{P}}$ and retrieval scheme $\Theta_{\mathbf{V}}$. Conditions under which the retrieval scheme is error-free and private have also been identified in Theorem 2. In this section, we will analyse the tradeoffs between the storage cost and the retrieval cost of our codes in the simplest special case when

$$\Phi = \{\{k\} : k \in \mathcal{K}\}.$$

Proposition 1: If $\mathbf{V}$ and $\mathbf{P}$ satisfy the retrievability condition in Theorem 2, then for any subset $\beta \subseteq \mathcal{K}$,

$$(T + L - R)(K - |\beta|) \leq \mathbf{rank}(\mathbf{P}(\beta))(T + L)$$

where $\mathbf{P}(\beta)$ is the submatrix of $\mathbf{P}$ by removing the rows indexed by $k \in \beta$.

Proof: Let β be a subset of $\mathcal{K}$. Consider again the system of equations in (28). Suppose the user is informed by a genie about the values of the following unknowns

$$\{(w_{t,k}, w^*_{\ell,k}) : t \in \mathcal{T}, \ell \in \mathcal{L} \text{ and } k \in \beta\}. \tag{45}$$

In that case, the number of remaining unknowns is now only $(T + L)(K - |\beta|)$. Given (45), the following equations

$$\sum_{t=1}^{T} v_{t,r,k} w_{t,k} + \sum_{\ell=1}^{L} v^*_{\ell,r,k} w^*_{\ell,k} = 0, \quad \forall k \in \beta, r \in \mathcal{R} \tag{46}$$

will become redundant. By direct counting, there are at most

$$\mathbf{rank}(\mathbf{P}(\beta))(T + L) + (K - |\beta|)R$$

non-redundant equations left in (28). Therefore, if the linear system (28) has a unique solution, then

$$(T + L)(K - |\beta|) \leq \mathbf{rank}(\mathbf{P}(\beta))(T + L) + (K - |\beta|)R, \tag{47}$$

or equivalently,

$$(T + L - R)(K - |\beta|) \leq \mathbf{rank}(\mathbf{P}(\beta))(T + L). \tag{48}$$

The proposition is thus proved. ∎

Proposition 2: If the matrix $\mathbf{V}$ satisfies the privacy condition (29), then $R \leq T$.

Proof: Consider any $k \in \mathcal{K}$ and the columns

$$\left[\begin{array}{c} v^*_{1,1,k} \\ \vdots \\ v^*_{L,1,k} \\ v_{1,1,k} \\ \vdots \\ v_{T,1,k} \end{array}\right] \ldots\ldots \left[\begin{array}{c} v^*_{1,R,k} \\ \vdots \\ v^*_{L,R,k} \\ v_{1,R,k} \\ \vdots \\ v_{T,R,k} \end{array}\right] \tag{49}$$

in $\mathbf{V}$. We may assume without loss of generality that all the columns are independent (otherwise, some of the columns can

be discarded). Now, if $R > T$, then

$$\mathbf{rank}(\langle V_{r,k}, r \in \mathcal{R}, k \in \alpha \rangle) > T.$$

As $\mathbf{rank}(\langle V_{r,k}, r \in \mathcal{R}, k \in \alpha, V_0 \rangle) \leq T + L$ and $\mathbf{rank}(V_0) = L$, we have $\mathbf{rank}(\langle V_{r,k}, r \in \mathcal{R}, k \in \alpha \rangle \cap V_0) > 0$. The proposition is thus proved. ∎

Theorem 3 (Tradeoff): For any storage code $\Delta_{\mathbf{P}}$ and retrieval scheme $\Theta_{\mathbf{V}}$ satisfying the retrievability and privacy condition,

$$1 \leq \mathsf{RC}(\Theta_{\mathbf{V}}, \Delta_{\mathbf{P}}) \left(K - \frac{1}{\mathsf{SC}(\Delta_{\mathbf{P}})} \right). \tag{50}$$

Equivalently,

$$\mathsf{RC}(\Theta_{\mathbf{V}}, \Delta_{\mathbf{P}}) \geq \frac{\mathsf{SC}(\Delta_{\mathbf{P}})}{K\mathsf{SC}(\Delta_{\mathbf{P}}) - 1}. \tag{51}$$

Proof: Let $\beta = \emptyset$. Applying Proposition 1, we have

$$(T + L - R)K \leq S(T + L) \tag{52}$$

$$T(K - S) \leq RK - L(K - S) \tag{53}$$

$$T \leq \frac{RK}{K - S} - L. \tag{54}$$

On the other hand, from Proposition 2, we have $R \leq T$. Therefore, we have

$$R \leq \frac{RK}{K - S} - L,$$

$$\frac{R}{L} \leq \frac{RK}{L(K - S)} - 1.$$

By (21) and (25), we have

$$1 \leq \frac{RK}{L(K - S)} - \frac{R}{L}$$

$$= \frac{R}{L(K - S)} \left(K - \frac{K - S}{1} \right).$$

Therefore,

$$1 \leq \mathsf{RC}(\Theta_{\mathbf{V}}) \left(K - \frac{1}{\mathsf{SC}(\Delta_{\mathbf{P}})} \right)$$

∎

Corollary 1: When $K = 3$, there does not exist any storage code $\Delta_{\mathbf{P}}$ and retrieval scheme $\Theta_{\mathbf{V}}$ satisfying the retrievability and privacy condition such that

$$\mathsf{RC}(\Theta_{\mathbf{V}}, \Delta_{\mathbf{P}}) = \mathsf{SC}(\Delta_{\mathbf{P}}) = 1/2.$$

Proof: A direct consequence of Theorem 3. ∎

Theorem 3 gives a tradeoff between storage cost and retrieval cost (in the context of our proposed classed of schemes). It has been well known that maximum distance separable (MDS) code has the lowest storage cost (compared to other storage codes with the same reliability). In the following theorem, we will show that the optimal storage-retrieval cost tradeoff can be achieved for MDS storage codes.

Theorem 4 (Optimality): Consider any fixed K, S. One can construct a $(K, K - S)$ maximum distance separable (MDS) code (specified by the parity check matrix $\mathbf{P}$) and an error-free, private retrieval scheme $\Theta_{\mathbf{V}}$ such that the lower bound (50) is tight. In other words,

$$\mathsf{RC}(\Theta_{\mathbf{V}}, \Delta_{\mathbf{P}}) = \frac{\mathsf{SC}(\Delta_{\mathbf{P}})}{K\mathsf{SC}(\Delta_{\mathbf{P}}) - 1}.$$

Proof: See Appendix B. ∎

V. Conclusion

Existing work on private information retrieval (PIR) problems largely focused on uncoded data storage (where every storage node stores all the data records). Their main focus was to design a retrieval scheme which has the lowest total upload (for transmitting retrieval queries) and download (for data retrieval) costs. One main drawback for uncoded storage is its high storage cost that the same data copy will be stored for multiple times. In this paper, we studied the problem of private information retrieval for coded data storage. Another major difference between this paper and existing works is that we focus only on the download cost. We have justified that the upload cost does not increase with the data record size. Therefore, in the scenario where data record size is large, the upload cost can become negligible. A fundamental question in PIR problems in coded storage is the characterisation of the tradeoffs between storage costs and retrieval costs. We have shown, to our surprise, that the tradeoff in fact depends on the number of data records in the system.

In the second half of the paper, we proposed a fairly general class of linear storage codes (where the underlying storage codes are specified by a parity check matrix) and an associated information retrieval scheme. We have identified conditions under which the retrieval scheme is error-free and private. One interesting property of our scheme is that the storage and retrieval costs are the same for any number of data records. Tradeoffs between the storage cost and the retrieval cost of our proposed schemes are then obtained. We also considered specific scenarios where MDS codes are used in the underlying coded storage. We have showed that a "randomly constructed" retrieval scheme is in fact optimal (in minimising the retrieval costs).

References

[1] A. G. Dimakis, P. B. Godfrey, Y. Wu, M. Wainwright, and K. Ramchandran, "Network Coding for Distributed Storage Systems," *Information Theory, IEEE Transactions*, vol. 56, pp. 4539 – 4551, 2010.

[2] S. Thakor, T. Chan, and K. W. Shum, "Symmetry in distributed storage systems," in *Proc. of 2013 IEEE International Symposium on Information Theory Proceedings (ISIT 2013)*, 2013.

[3] M. A. Tebbi, T. Chan, and C. W. Sung, "Linear programming bounds for robust locally repairable storage codes," in *to appear in Information Theory Workshop*, 2014.

[4] K. V. Rashmi, N. B. Shah, P. V. Kumar, and K. Ramchandran, "Explicit construction of optimal exact regenerating codes for distributed storage," *47th Annual Allerton Conference on Communication Control and Computing Allerton*, 2009.

[5] N. B. Shah, K. V. Rashmi, P. V. Kumar, and K. Ramchandran, "Explicit codes minimizing repair bandwidth for distributed storage," *IEEE Information Theory Workshop 2010 ITW 2010*, p. 11, 2009. [Online]. Available: http://arxiv.org/abs/0908.2984

[6] B. Chor, O. Goldreich, E. Kushilevitz, and M. Sudan, "Private information retrieval," in *Proceedings of the 36th Annual Foundations of Computer Science*, 1995, pp. 41–50.

[7] B. Chor, E. Kushilevitz, O. Goldreich, and M. Sudan, "Private information retrieval," *Journal of the ACM (JACM)*, vol. 45, no. 6, pp. 965–981, 1998.

[8] E. Kushilevitz and R. Ostrovsky, "Replication is not needed: Single database, computationally-private information retrieval," in *FOCS*, 1997, pp. 364–373.

[9] S. Yekhanin, "Locally decodable codes and private information retrieval schemes," Ph.D. dissertation, Cambridge, MA, USA, 2007, aAI0819886.

[10] C. Gentry and Z. Ramzan, "Single-database private information retrieval with constant communication rate," in *ICALP*, 2005, pp. 803 – 815.

[11] H. Lipmaa, "First CPIR protocol with data-dependent computation," in *ICISC*, 2009, pp. 193–210.

[12] A. Beimel and Y. Stahl, "Robust information-theoretic private information retrieval," in *Proceedings of the 3rd International Conference on Security in Communication Networks (SCN'02)*, 2002, pp. 326–341.

[13] C. Devet, I. Goldberg, and N. Heninger, "Optimally robust private information retrieval," in *Proceedings of the 21st USENIX Conference on Security Symposium*, ser. Security'12. Berkeley, CA, USA: USENIX Association, 2012, pp. 13–13.

[14] N. B. Shah, K. V. Rashmi, and K. Ramchandran, "One extra bit of download ensures perfectly private information retrieval," in *Information Theory Proceedings (ISIT), 2014 IEEE International Symposium on*, (Hawaii, USA) 2014.

[15] D. S. Papailiopoulos and A. G. Dimakis, "Locally repairable codes," in *Information Theory Proceedings (ISIT), 2012 IEEE International Symposium on*. IEEE, 2012, pp. 2771–2775.

[16] S.-W. Ho, C. W. Tan, and R. Yeung, "Proving and disproving information inequalities," in *2014 IEEE International Symposium on Information Theory (ISIT)*, June 2014, pp. 2814–2818.

[17] J. T. Schwartz, "Fast probabilistic algorithms for verification of polynomial identities," *Journal of the ACM (JACM)*, vol. 27, no. 4, pp. 701–717, 1980.

[18] R. Zippel, *Probabilistic algorithms for sparse polynomials*. Springer, 1979.

APPENDIX A
PROOF OF THEOREM 1

Suppose to the contrary that the storage-retrieval cost tuple $(1/2, 1/2)$ is strongly achievable. By definition, there exists a sequence of $(N^{(a)}, 3)$ storage codes $\Omega^{(a)}$ and Φ-private, error-free retrieval schemes $\Lambda^{(a)} = (\Lambda_1^{(a)}, \Lambda_{2,k}^{(a)})$ such that

$$\limsup_{a\to\infty} \mathsf{SC}(\Omega^{(a)}) \le \frac{1}{2} \tag{55}$$

$$\limsup_{a\to\infty} \mathsf{RC}(\Lambda^{(a)}, \Omega^{(a)}) \le \frac{1}{2} \tag{56}$$

$$\lim_{a\to\infty} N^{(a)} = \infty. \tag{57}$$

For each a, let $(D_n^{(a)}, n = 1, \ldots, N^{(a)})$ be the data records and $(X_k^{(a)}, k = 1, \ldots, K^{(a)})$ be the content stored at the storage nodes. The retrieval scheme is defined by $(M^{(a)}, Q_k^{(a)}, A_k^{(a)}, k \in \mathcal{K})$. For notation simplicity, we will often drop the superscript $^{(a)}$ in the remainder of the proof.

First, notice that $(D_n, n \in \mathcal{N})$ is mutually independent (see also (7)). Therefore, for any $k \in \mathcal{K}$,

$$I(Q_k A_k; D_{\mathcal{N}}) \ge \sum_{n\in\mathcal{N}} I(Q_k A_k; D_n). \tag{58}$$

On the other hand,

$$I(Q_k A_k; D_{\mathcal{N}}) = I(Q_k; D_{\mathcal{N}}) + I(A_k; D_{\mathcal{N}}|Q_k) \tag{59}$$

$$\le H(A_k|Q_k). \tag{60}$$

Here, the last inequality follows from (5). Therefore

$$H(A_k|Q_k) \ge \sum_{n\in\mathcal{N}} I(Q_k A_k; D_n) \tag{61}$$

and consequently,

$$\limsup_{a\to\infty} \frac{\max_{k\in\mathcal{K}} \sum_{n=1}^{N} I(Q_k A_k; D_n)}{H(D_1)}$$

$$\le \limsup_{a\to\infty} \frac{\max_{k\in\mathcal{K}} H(A_k|Q_k)}{H(D_1)} \tag{62}$$

$$\le \frac{1}{2} \tag{63}$$

where the last inequality follows from (55). Hence, for any $c > 1/2$ and sufficiently large a, we have

$$c \ge \frac{\max_{k\in\mathcal{K}} \sum_{n=1}^{N} I(Q_k A_k; D_n)}{H(D_1)}. \tag{64}$$

Consequently,

$$\frac{\sum_{k\in\mathcal{K}} \sum_{n=1}^{N} I(Q_k A_k; D_n)}{N} \le \frac{cKH(D_1)}{N}. \tag{65}$$

Therefore, there must exist at least one n^* such that

$$\sum_{k\in\mathcal{K}} I(Q_k A_k; D_{n^*}) \le \frac{cKH(D_1)}{N}$$

which implies that for all $k \in \mathcal{K}$,

$$I(Q_k A_k; D_{n^*}) \le \frac{cKH(D_1)}{N}. \tag{66}$$

Now, notice that

$$\begin{aligned} &H(M, Q_k, X_k, D_{n^*}, A_k) \\ &\quad = H(M) + H(Q_k|M) + H(X_k D_{n^*}|M, Q_k) \\ &\qquad + H(A_k|M, Q_k, X_k, D_{n^*}) \\ &\quad \overset{(i)}{=} H(M) + H(Q_k) + H(X_k D_{n^*}|Q_k) \\ &\qquad + H(A_k|Q_k, X_k, D_{n^*}) \\ &\quad = H(M) + H(Q_k, X_k, D_{n^*}, A_k) \end{aligned}$$

where (i) follows from (5), (6) and (9). Hence, M and (Q_k, X_k, D_{n^*}, A_k) are independent. Consequently, for sufficiently large a,

$$I(D_{n^*}; Q_k A_k|M = n^*) = I(D_{n^*}; Q_k A_k) \tag{67}$$

$$\le \frac{cKH(D_1)}{N} \tag{68}$$

$$\overset{(ii)}{=} \frac{cKH(D_{n^*})}{N}. \tag{69}$$

Here, (ii) follows from that $H(D_n)$ is constant for all n. Similarly, by using (56) and the privacy condition (9), we can prove that for sufficiently large a,

$$H(A_k|Q_k, M = n^*) \le cH(D_{n^*}). \tag{70}$$

In addition, by (5), (6) and (8), we have

$$I(X_{\mathcal{K}}, D_{n^*}; Q_{\mathcal{K}}|M=n^*)=0 \tag{71}$$
$$H(D_{n^*}|Q_{\mathcal{K}}, A_{\mathcal{K}}, M=n^*)=0 \tag{72}$$
$$H(A_k|Q_k, X_k, M=n^*)=0 \tag{73}$$

Before we continue, we will need the following intermediate lemma.

Lemma 1:

$$\begin{aligned}&2H(D_{n^*})+H(A_1|Q_{\mathcal{K}}A_2D_{n^*})\\&\quad\le 2H(A_1|Q_1)+H(A_2|Q_2X_2)+2H(A_3|Q_3)\\&\quad+H(D_{n^*}|Q_{\mathcal{K}}A_{\mathcal{K}})+I(Q_2A_2;D_{n^*})+I(Q_{\mathcal{K}};X_2D_{n^*}).\end{aligned} \tag{74}$$

Consequently,

$$\begin{aligned}&2H(D_{n^*}|M=n^*)+H(A_1|Q_{\mathcal{K}}A_2D_{n^*},M=n^*)\\\le&2H(A_1|Q_1,M=n^*)+H(A_2|Q_2X_2,M=n^*)\\&+2H(A_3|Q_3,M=n^*)+H(D_{n^*}|Q_{\mathcal{K}}A_{\mathcal{K}},M=n^*)\\&+I(Q_2A_2;D_{n^*}|M=n^*)+I(Q_{\mathcal{K}};X_2D_{n^*}|M=n^*).\end{aligned} \tag{75}$$

Proof: Using [16], we can obtain the following equality which can be directly verified by re-writing each term in terms of joint entropy.

$$\begin{aligned}&-2H(D_{n^*})-H(A_1|Q_{\mathcal{K}}A_2D_{n^*})\\=&H(A_2|Q_{\mathcal{K}},D_{n^*},A_1,A_3,X_2)+I(A_1;Q_3|Q_1,Q_2,A_2)\\&+H(A_3|Q_{\mathcal{K}}A_1A_2D_{n^*}X_2)+I(A_1;A_2|Q_{\mathcal{K}}D_{n^*}X_2)\\&+I(A_1;Q_2|Q_1A_2)+I(A_1;A_2|Q_1)+I(A_1;A_3|Q_{\mathcal{K}}A_2)\\&+I(Q_1;A_2|Q_2Q_3)+I(Q_2;D_{n^*}|X_2)+I(Q_2;X_2)\\&+I(Q_1;A_3|Q_3A_2)+I(A_2;A_3|Q_3)+I(Q_3;A_2|Q_2)\\&+I(Q_2;A_3|Q_1Q_3A_2)+I(Q_3;X_2|Q_1Q_2A_2D_{n^*})\\&+I(A_2;D_{n^*}|Q_2X_2)+I(A_2;A_3|Q_{\mathcal{K}}A_1X_2D_{n^*})\\&+I(Q_1;X_2|Q_2A_2D_{n^*})+I(X_2;A_3|Q_{\mathcal{K}}A_1A_2D_{n^*})\\&-2H(A_1|Q_1)-H(A_2|Q_2X_2)-2H(A_3|Q_3)\\&-H(D_{n^*}|Q_{\mathcal{K}}A_{\mathcal{K}})-I(Q_2A_2;D_{n^*})-I(Q_{\mathcal{K}};X_2D_{n^*}).\end{aligned}$$

Hence, the lemma follows from the nonnegativity of mutual information and entropies. ∎

Now, by invoking Lemma 1, and (69) – (73), we have

$$H(A_1|Q_{\mathcal{K}}A_2D_{n^*},M=n^*)\le\left(4c-2+\frac{cK}{N}\right)H(D_{n^*}).$$

Similarly, for $i\neq j$,

$$H(A_i|Q_{\mathcal{K}}A_jD_{n^*},M=n^*)\le\left(4c-2+\frac{cK}{N}\right)H(D_{n^*})$$

by symmetry. Finally, for any $n\in\mathcal{N}$,

$$\begin{aligned}&H(D_n|X_1,D_{n^*})\\&\overset{(iii)}{=}H(D_n|X_1,D_{n^*},Q_{\mathcal{K}},M=n) &(76)\\&\le H(D_n,A_{\mathcal{K}}|X_1,D_{n^*},Q_{\mathcal{K}},M=n) &(77)\\&\overset{(iv)}{=}H(A_{\mathcal{K}}|X_1,D_{n^*},Q_{\mathcal{K}},M=n) &(78)\\&\le\sum_{k\neq1}H(A_k|X_1,A_1,D_{n^*},Q_{\mathcal{K}},M=n)\\&\qquad+H(A_1|X_1,D_{n^*},Q_{\mathcal{K}},M=n) &(79)\\&=\sum_{k\neq1}H(A_k|X_1,A_1,D_{n^*},Q_{\mathcal{K}},M=n)\\&\le K\left(4c-2+\frac{cK}{N}\right)H(D_{n^*}) &(80)\end{aligned}$$

where (iii) follows from (5) and (iv) from (8). Hence,

$$\begin{aligned}H(D_{\mathcal{N}}|X_1,D_{n^*})&\le\sum_{n\in\mathcal{N}}H(D_n|X_1,D_{n^*}) &(81)\\&\le NK\left(4c-2+\frac{cK}{N}\right)H(D_{n^*}). &(82)\end{aligned}$$

Consequently,

$$\begin{aligned}&NH(D_{n^*})\\&=H(D_{\mathcal{N}})\\&\le H(X_1,D_{n^*})+H(D_{\mathcal{N}}|X_1,D_{n^*})\\&\le H(X_1)+H(D_{n^*})+H(D_{\mathcal{N}}|X_1,D_{n^*})\\&\le H(X_1)+H(D_{n^*})+NK\left(4c-2+\frac{cK}{N}\right)H(D_{n^*}).\end{aligned}$$

Therefore,

$$1\le\frac{H(X_1)}{NH(D_{n^*})}+\frac{1}{N}+K\left(4c-2+\frac{cK}{N}\right). \tag{83}$$

Consequently,

$$\liminf_{a\to\infty}\frac{H(X_1)}{NH(D_{n^*})}+\frac{1}{N}+K\left(4c-2+\frac{cK}{N}\right)\ge1.$$

As c is arbitrarily close to $1/2$, together with (57), we have

$$\liminf_{a\to\infty}\frac{H(X_1)}{NH(D_{n^*})}\ge1.$$

In other words, $\mathsf{SC}(\Omega^{(a)})\ge1$. The theorem is proved.

Appendix B
Proof of Theorem 4

Consider any K and $S<K$. For sufficiently large field $GF(q)$, we can always construct a $(K,K-S)$ MDS code. Let its $K\times S$ parity check matrix be $\mathbf{P}$ and $\Delta_{\mathbf{P}}$ be the corresponding induced storage code. Hence,

$$\mathsf{SC}(\Delta_{\mathbf{P}})=\frac{1}{K-S}.$$

Let $R=T=K-S$ and $L=S$. Hence, $T+L=K$ and the matrix $\mathbf{V}$ is indeed a square matrix. Now consider the

retrieval scheme $\Theta_{\mathbf{V}}$. It can be verified easily that

$$\mathsf{RC}(\Theta_{\mathbf{P}}, \Delta_{\mathbf{V}}) = \frac{1}{S}$$

and hence the equality in (50) holds.

To prove the theorem, it suffices to prove that one can construct a matrix $\mathbf{V}$ such that its induced retrieval scheme $\Theta_{\mathbf{V}}$ is error-free and private. First of all, for notation simplicity, we rewrite the variables in (28) as follows:

$$z_{i,k} = \begin{cases} w^*_{i,k} & \text{if } 1 \le i \le L \\ w_{i-L,k} & \text{if } 1 \le i - L \le T. \end{cases} \tag{84}$$

$$y_{i,r,k} = \begin{cases} v^*_{i,r,k} & \text{if } 1 \le i \le L \\ v_{i-L,r,k} & \text{if } 1 \le i - L \le T \end{cases} \tag{85}$$

Using the new notation, (28) becomes

$$\sum_{k=1}^{K} p_{k,s} z_{i,k} = 0, \quad \forall i \in \mathcal{K}, s \in \mathcal{S} \tag{86}$$

$$\sum_{i=1}^{K} y_{i,r,k} z_{i,k} = 0, \quad \forall k \in \mathcal{K}, r \in \mathcal{R}. \tag{87}$$

As $\mathbf{P}$ is the $K \times S$ parity-check matrix of a $(K, K-S)$ MDS code. Hence, for any $i \in \mathcal{K}$ and a subset $\beta \subseteq \mathcal{K}$ of size at least $K - S$, if

$$\sum_{k=1}^{K} p_{k,s} z_{i,k} = 0, \quad \forall s \in \mathcal{S}$$

and $z_{i,k} = 0$ for all $k \in \beta$, then $z_{i,k} = 0$ for all $k \in \mathcal{K}$.

Next, let

$$y_{i,r,k} \triangleq \begin{cases} 1 & \text{if } 0 = k - i - r \mod K \\ 0 & \text{otherwise} \end{cases} \tag{88}$$

for $i, k \in \mathcal{K}$ and $r \in \mathcal{R}$. Let

$$\square \triangleq \{(i,k) : \exists r \in \mathcal{R} \text{ such that } 0 = k - i - r \mod K\}.$$

Then (87) becomes

$$z_{i,k} = 0, \quad \forall (i,k) \in \square.$$

Also, for any fixed $i \in \mathcal{K}$, the set $\{k : (i,k) \in \square\}$ has exactly $R = K - S$ elements. Therefore, by the construction of $\mathbf{P}$, we can conclude that $z_{i,k} = 0$ for all $i \in \mathcal{K}$. Hence, for our choice of $y_{i,r,k}$ (and hence, $v^*_{i,r,k}$ and $v_{i,r,k}$), the retrievability condition is satisfied.

So far, we proved the existence of $\mathbf{V}$ (via (85) and (88)) such that the retrievability condition will be satisfied. However, the chosen $\mathbf{V}$ may not satisfy the privacy condition (29). In the following, we will prove the existence of a matrix $\mathbf{V}$ which can satisfy both conditions.

The proof for existence relies much on the well-known Schwartz-Zippel Lemma [17], [18] which is restated as follows.

Lemma 2 (Schwartz-Zippel Lemma [17], [18]): Let $F[b_1, \ldots, b_n]$ be a non-zero polynomial with degree $\gamma \ge 0$ over the finite field $GF(q)$. If $b_1, \ldots, b_n$ are randomly and independently chosen from the field, then the probability that $F[b_1, \ldots, b_n] = 0$ will be at most d/q.

Now, consider again the system of linear equations (28). The retrievability condition will be satisfied if and only if the determinant of the coefficient matrix for (28) is non-zero. Let

$$F(v^*_{\ell,r,k}, v_{t,r,k}, t \in \mathcal{T}, r \in \mathcal{R}, k \in \mathcal{K}, \ell \in \mathcal{L})$$

be the determinant function. Here, we assume that $\mathbf{P}$ is given and fixed. Clearly, the determinant function F is a non-zero polynomial because we have already shown earlier that it has non-zero values for some choices of

$$v^*_{\ell,r,k}, v_{t,r,k}, t \in \mathcal{T}, r \in \mathcal{R}, k \in \mathcal{K}, \ell \in \mathcal{L}.$$

Applying the Schwartz-Zippel Lemma, for sufficiently large q, the determinant function is non-zero with arbitrarily high probability. In other words, the retrievability condition holds with high probability, if

$$v^*_{\ell,r,k}, v_{t,r,k}, t \in \mathcal{T}, r \in \mathcal{R}, k \in \mathcal{K}, \ell \in \mathcal{L}$$

are in fact chosen randomly (and when q is sufficiently large). Finally, to prove the theorem, it suffices to prove that the privacy condition will also be satisfied by a randomly generated $\mathbf{V}$ with arbitrarily high probability.

Notice the columns

$$V_{r,k} = [v^*_{1,r,k}, \ldots, v^*_{L,r,k}, v_{1,r,k}, \ldots, v_{T,r,k}]^\top.$$

Since each element $v^*_{\ell,r,k}$ or $v_{t,r,k}$ is randomly and independently selected and $\mathbf{rank}(V_0) = S$, the probability that

$$\langle V_{r,k}, r = 1, \ldots R \rangle \cap V_0 = \{0\} \tag{89}$$

is equal to

$$\prod_{c=0}^{K-S-1} \left(1 - \frac{q^{S+c}}{q^K}\right) \ge \left(1 - \frac{1}{q}\right)^{K-S}.$$

Hence, the probability can be made arbitrarily close to 1 for sufficiently large q. In other words, the privacy condition will also be satisfied by a randomly generated $\mathbf{V}$. The theorem is proved.